

 APPLICATION TARGET: SENIOR APPLIED AI & ML ENGINEER – HEALTH AI & CLINICAL DECISION SUPPORT

TARGETED EXECUTIVE SUMMARY

Senior Applied AI & Machine Learning Engineer with 7+ years of experience building **calibrated predictive models, clinical decision support systems, and human-in-the-loop AI workflows**. Specialized in bridging raw biometric and health telemetry data with empathetic, reliable AI recommendations. Ranked in the **Top 1.2% worldwide on Kaggle (Expert, #734 of 61,348+)**. Creator of **Clinical Risk Prediction Pipeline** (15 model iterations optimized for Brier Score calibration) and **MyHealth Decoder** (built for the Google Med-Gemma Impact Challenge). Passionate about empowering human coaches and members with safe, actionable, and non-hallucinating health intelligence.

SPECIALIZED TECHNICAL STACK

- **Clinical & Health AI:** MedGemma, BioMistral, Clinical Risk Calibration, Brier Score Minimization, Isotonic Regression, Medical Terminology Extraction, Human-in-the-Loop Safeguards.
- **Predictive & Tabular ML:** LightGBM, CatBoost, XGBoost, Scikit-Learn, Feature Importance (SHAP/Permutation), Time-Series Biometric Tracking.
- **Computer Vision & Medical Imaging:** PyTorch, Torchvision, Multi-Model Ensembling, Triple-Fusion Blends (RSNA Knee Abnormality 0.9576 Apex score).
- **AI Pipelines & Cloud Systems:** FastAPI, Python, Docker, ChromaDB, Cloudflare Workers, PostgreSQL, SQLite.

KEY FEATURED ENGINEERING & PROJECTS

Clinical Risk Prediction Pipeline (Cardiovascular Decision Support)

Machine Learning Architect | Portfolio: kiza-portfolio.pages.dev

- Architected an end-to-end clinical risk prediction pipeline evaluating 13 multi-modal patient features.
- Evaluated **15 model architectures** to establish the optimal balance between discrimination (F1 score) and **probability calibration (Brier Score)**.
- Ensured predictions reflect true biological risk probabilities rather than overconfident binary classifications, setting a benchmark for safe clinical decision support.

MyHealth Decoder — MedGemma Patient Advocate

AI Lead | Hugging Face: [kizabgd123/myhealth-decoder](https://huggingface.co/kizabgd123/myhealth-decoder)

- Engineered for the **Google Med-Gemma Impact Challenge** to translate opaque diagnostic reports, lab results, and physician notes into patient-centric, empathetic summaries.
- Implemented strict multi-turn safety guardrails to ensure outputs never prescribe unverified interventions while reinforcing trust and adherence.

RSNA Knee Abnormality Detection (Medical Imaging SOTA)

Competition Engineer | Kaggle Kernel

- Built a deep learning computer vision pipeline analyzing complex knee MRI scans for structural abnormalities.
- Developed a **triple-model fusion ensemble** achieving a competitive **0.9576 Apex AUC score**, demonstrating top-tier competence in handling noisy, real-world clinical datasets.

JudgeGuard Deterministic Reality Verifier

Creator | GitHub: [kizabgd123/judge-guard-core-master](https://github.com/kizabgd123/judge-guard-core-master)

- Designed a zero-trust governance runtime that eliminates hallucinations in LLM workflows.
- Applied strict semantic drift monitoring and Layer 4 physical outcome verifications to ensure that all recommendations sent to users are grounded in verified telemetry.

PROFESSIONAL EXPERIENCE

Senior Applied AI Engineer (Health & Predictive Systems) | Independent Consultant

2023 – Present

- Designed personalized recommendation engines and behavior-tracking pipelines combining user feedback with calibrated ML models.
- Implemented RAG knowledge bases connecting clinical literature and wellness guidelines to real-time coach copilots.
- Built automated data validation gates ensuring biometric stream integrity prior to model scoring.

Lead Machine Learning & Analytics Engineer | Data & Core Systems

2020 – 2023

- Led development of predictive modeling pipelines on high-stakes tabular datasets.
- Established company-wide model governance, calibration audits, and drift detection pipelines.
- Integrated automated CI/CD and monitoring using Datadog Synthetics to track inference latency and prediction reliability.

HONORS, AWARDS & IMPACT

- **Kaggle Notebooks Expert**: Ranked **#734 / 61,348+** worldwide (Top 1.2%).
- **Google Med-Gemma Impact Challenge Participant**.
- **GitHub Developer Program Member**.